Final

IN-FLIGHT PERFORMANCE EVALUATION OF THE DEEP SPACE 1 AUTONOMOUS NAVIGATION SYSTEM

Shyam BHASKARAN, Joseph E. RIEDEL, Stephen P. SYNNOTT, T. C. WANG, Robert A. WERNER, Brian M. KENNEDY

Navigation and Flight Mechanics Section
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
USA

Deep Space 1 is the first mission in NASA's New Millennium Program, a series of missions to demonstrate the feasibility of new technologies for spaceflight. DS1 was launched on October 24, 1999 and has completed its first leg of its mission -- flyby of the asteroid Braille -- on July 29, 1999. An additional encounter is planned with the short period comet Borrelly in September 2001. The new technologies being demonstrated on DS1 include an ion propulsion system to provide maneuvering thrust, a combined visible/infrared/ultraviolet imaging instrument named MICAS (Miniature Integrated Camera and Spectrometer), and an autonomous navigation system. This paper describes the computational elements of the autonomous navigation system and assesses its performance in guiding the spacecraft to its first target. Some of the difficulties encountered, and how they were overcome, will also be described.

1 – INTRODUCTION

Standard navigation techniques for interplanetary spacecraft involve a combination of radio (two-way coherent Doppler and ranging) data, obtained by tracking the spacecraft using antennas at JPL's Deep Space Network (DSN) tracking stations, augmented by optical data from an onboard camera during encounters. This combination is very accurate, and has been used successfully to navigate to all planets except Pluto, several asteroids, and comet Halley. However, in order to fully realize NASA's vision of the future of deep-space exploration with multiple small, inexpensive spacecraft roaming the solar system, it is desirable to automate the processes required for interplanetary missions, including navigation. It has been known for many years that the navigation system could be fully automated and self-contained by using a camera taking triangulation images of solar system bodies to determine the spacecraft's position and velocity, and then computing and executing maneuvers onboard to deliver the spacecraft to its target. Such a system has been developed for JPL's Deep Space 1 (DS1) asteroid/comet flyby mission. Although the principles behind this system are fairly simple, the implementation posed several challenges when applied to the flight of an actual spacecraft. The purpose of this paper is twofold; first to briefly describe the computational elements of the navigation system, and second to assess its performance during the first leg of the DS1 mission.

2 – THE MISSION

DS1 is the first of the interplanetary missions of NASA's New Millennium Program (a series of missions specifically chosen to test new and untried technologies for spaceflight). In addition to autonomous navigation, other primary technologies being demonstrated include the first use of an ion propulsion system for trajectory control, an advanced solar array for

power, and a low-mass imaging system named MICAS (Miniature Integrated Camera and Spectrometer), (see [Raym, 96] for a complete description of all the new technologies on DS1).

DS1 was launched on a Delta 7326 rocket on October 24, 1998, and flew by its first target, the asteroid 9969 Braille, on July 29, 1999. Although originally planned to flyby two more targets, the loss of the onboard star tracker in November, 1999 prevented the spacecraft from performing its nominal thrust profile to provide the energy needed to reach both of these targets. As of this writing, the spacecraft team is developing software to use the MICAS imaging camera to replace the star tracker. The goal is to start another thrust profile in the summer of 2000 to flyby one of the targets, the comet Borrelly, in September 2001.

From an autonomous navigation (autonav) standpoint, the two other technologies having the largest impact are the use of ion propulsion engines, and the characteristics of the MICAS camera. Unlike chemical propulsion systems which burn for short periods of time at very high thrust, the IPS produces very little thrust but is capable of operating for very long periods of time. Ionized xenon is accelerated by passing it through an electrically charged grid before exiting the nozzle. The resulting thrust is on the order of millinewtons, with specific impulses reaching values in the thousands of seconds (as compared to 200-400 seconds for chemical rockets). The thrust can be throttled by varying the voltage on the grids; for DS1, the IPS has about 100 throttle levels, with a thrust range of 20 to 90 mN. Since electrical power is generated from the solar arrays, the maximum achievable thrust depends on the distance to the sun. IPS trajectories are characterized by long thrusting periods of weeks to months, interspersed with coast arcs when the IPS is shut off. For DS1, the thrusting periods have the dual purpose of providing enough energy to the spacecraft to reach its targets, and correcting launch injection, orbit determination, and maneuver execution errors to achieve the desired targeting conditions.

The MICAS camera has four channels, two in the visible light spectrum, and an infrared and an ultraviolet spectrometer; only the visible light channels were used by autonav. Of these, the one used during the majority of cruise and approach to the asteroid was a standard Charge-Coupled-Device (CCD) chip with a 1024 square pixel array. Each pixel has a field-of-view (FOV) of about 13 μrad for a total FOV in the CCD of 1.3 mrad, or 0.76 deg. The other visible channel is an experimental Active Pixel Sensor (APS) array which has a 256 square pixel array. Each pixel of the APS has a FOV of 18 μrad for a total FOV of 4.6 mrad (0.26 deg). Both were coupled to a telescope with a focal length of 685 mm whose boresight is fixed to the spacecraft. Also, both have 12 bit digitization, resulting in data numbers (DN) values for each pixel ranging between 0 (no signal) and 4095 (saturation). Prior to launch, the MICAS development team had determined that the CCD would have excessive charge bleeding when exposed to a bright, extended object such as during the flyby period (MICAS has no shutter). Thus, it was decided to use the far less sensitive APS in the final 20 minutes or so of terminal tracking during flyby. This decision had a major impact on the success of the flyby tracking, as will be described later.

3 – COMPUTATIONAL ELEMENTS OF THE AUTONAV SYSTEM

The entire autonav code was designed to be as self-contained and modular as possible to make it adaptable to other missions. The system can be divided into three components; an executive which is responsible for scheduling and executing events, a real-time ephemeris

server which provides spacecraft and solar system body ephemeris information to other flight software elements, and the computational elements which perform the fundamental navigation updates. The first two components will not be described; for more information, see [Ried 97]. This paper is primarily concerned with the performance of the computational elements, which include the following functions: 1) orbit determination, 2) maneuver planning, and 3) encounter target tracking. Each of these will now be briefly described.

3.1 – Orbit Determination

Orbit determination (OD) is the process by which the spacecraft's state (position and velocity) and other parameters relevant to the trajectory, such as nongravitational accelerations acting on the spacecraft, are estimated. In order to keep this process as self-contained onboard the spacecraft as possible, the only data used by autonav to obtain an OD solution are images of solar system bodies (asteroids in this case) taken by the MICAS camera. The principle is the following; each sighting of an asteroid in the camera FOV places the spacecraft in a position along that line-of-sight (LOS). Two or more such sightings of different asteroids fixes the spacecraft's three-dimensional position by triangulation. The stars in the background are needed to determines the inertial pointing direction of the camera boresight (since the stars are so distant, their inertial directions will not change measurably when seen from different locations in the solar system, so they can be thought of as "fixed" in the sky). In practice, however, two simultaneous sightings are not practical with one camera, and instead, a series of LOS fixes are taken of several asteroids. Several clusters of sightings are then incorporated into a least-squares filter to obtain an OD solution. The accuracy of this type of data is dependent on several factors, but previous analysis has shown that it is capable of meeting the navigation requirements of asteroid flyby missions. For clarity in the following description, the term "beacons" refers to the asteroids used for navigation, whereas the target is the object being encountered (asteroid Braille for the first leg of the mission).

3.1.1 - Image Processing

The purpose of image processing is to predict the locations of beacons and surrounding stars at given times, determine the center of the asteroid and the stars in the camera frame, and compute the associated pointing of the camera boresight. The ability of the navigation system to perform autonomously hinges on its ability to accurately perform the centerfinding and ensuring that bad data do not corrupt the OD solution.

Predicting the location of beacon asteroids is the simplest of these procedures. A list of beacon asteroids to observe is stored onboard the spacecraft, along with ephemerides of all the beacons. At predetermined times, the nominal spacecraft trajectory is differenced with the ephemeris of a given beacon to get the relative pointing vector. This information is then passed to the spacecraft attitude control system (ACS) which slews the spacecraft to the correct attitude and shutters the picture with the provided exposure length. After the frame is taken, it is stored as a file in the main flight computer.

The image processing itself has two stages; the first performs a coarse registration of the beacon and stars in the frame, and the second fine tunes the center locations to subpixel accuracy. The second step uses a cross-correlation technique inherited from the Galileo mission (for a detailed mathematical description of this process, see [Vaug 92]. A post-processor then screens the results for low signal and bad data by deleting objects which do not

pass a threshold. From prior experience with the Galileo mission and from tests done on images taken from a ground telescope, pre-flight estimates on the accuracy of this procedure was around 0.1 pixels. For reasons described later, the actual performance in flight varied between 0.2 and 0.8 pixels.

3.1.2 - Dynamic Equations

The spacecraft trajectory dynamic model employed by autonav includes, in addition to the central body acceleration, third body perturbations from other planets, solar radiation pressure, thrust forces from the IPS, and a general bias acceleration term used for small unmodelled forces. The propagation of the spacecraft motion is done entirely in a heliocentric, Earth Mean Equator of 2000 coordinate system using a 7-8th order Runge-Kutta numerical integrator. Third body perturbations included those from Venus, Earth, Moon, Mars, Jupiter, and Saturn. The model of solar radiation pressure assumed a spherical shape for the spacecraft whose area was equal to the area of the panels. The unmodelled acceleration term was a bias in three cartesian axes applied over the whole data arc. Nominally, this term is set to zero but is estimated in the filter to absorb residual accelerations.

3.1.3 - Filter

The autonav system uses the standard practice of linearizing about a nominal trajectory and estimating corrections to the nominal state parameters [Lieb 67] which minimize the data residuals in a least-squares sense. The estimated parameters include corrections to the spacecraft state (position and velocity), the bias accelerations, and the IPS thrust scale factors. The batch epoch-state filter is used whereby all the data in a given batch are used to estimate the state at the epoch time of the batch, after which the solution and associated covariance are propagated through the data arc to the current time. The batch length for computing solutions varied, but in general, was about 30 days, which was found from preflight studies to be an optimal length.

A brief description of the OD process is as follows. At roughly weekly intervals, the spacecraft sets aside a 4-6 hour block of time for beacon asteroid observations. Using a list of pre-planned beacons and its stored spacecraft and asteroid ephemerides, the spacecraft turns and shoots a series of frames of each beacon. Each exposed frame is immediately processed by the image processing link, and the pertinent information appended to a file. After all the frames were taken and processed, the OD link computes updates to the trajectory using all the data in the arc; the solution is then mapped to the current time and into the future. A new spacecraft ephemeris file is prepared, and becomes the nominal one until the next OD solution. Thus, the spacecraft has an updated knowledge of its own whereabouts. Various cleanup tasks are also performed, such as truncating the history and/or image data files. Finally, if called for, a maneuver targeting computation is performed, which will now be described.

3.2 - Maneuver Planning

The purpose of the maneuver planner is to compute the course corrections needed to achieve the target flyby conditions. The course correction can be implemented in one of three ways: changing the magnitude, direction, and duration of the nominal thrust profile that the ion

engine flies, adding a Trajectory Correction Maneuver (TCM) using the ion engine, or adding a TCM using the hydrazine attitude control thrusters. In all three cases, the correction maneuver is computed by taking partial derivatives of the desired target condition with respect to the control parameters, inverting this matrix, and multiplying by the residual error formed by the difference between the current and desired target condition. Thus, it is a linear control process and requires a reasonably good starting condition in order to converge. The initial design of the nominal trajectory is done on the ground and uplinked to the spacecraft. More details of the maneuver design process can be found in [Desa 97] and [Ried 97].

3.3 - Encounter Target Tracking

The complete set of dynamics and filter described above was used as the primary mode of operation for OD solutions throughout cruise and until several hours prior to encounter. At this stage, no more maneuvers are performed and the main purpose of autonav is to maintain visual lock on the target as it flies by. This requires rapid updates of the state as tracking images are taken, and the OD link is not fast enough for this purpose. For this reason, a compact filter, termed Reduced State Encounter Navigation (RSEN) was developed. RSEN uses the final position and velocity from the main OD link at 30 minutes prior to encounter and models the trajectory as a target-centered straight line through encounter. As frames are taken and processed, the epoch position alone is updated and a new linear course computed. The spacecraft relative target position and updated time of closest approach are passed to the ACS system to keep the target in the camera FOV. An important difference between RSEN and the mainline OD is that it uses the smaller FOV and less sensitive APS channel; no CCD images are processed by RSEN. For more information on the details of RSEN, see [Bhas 98].

4 - OPERATIONAL RESULTS

4.1 - MICAS Problems

Prior to launch, all the computational elements (except RSEN) of the autonav system were extensively tested by Monte Carlo simulations using realistic models for the performance of the engine and the spacecraft ACS. These tests indicated that autonav was capable of delivering the spacecraft to its close flyby of an asteroid without need of ground intervention and within reasonable margins for propellant usage [Bhas 98]. The actual flight, however, turned up several unanticipated problems, largely due to the performance of the MICAS camera, which made autonav extremely difficult to accomplish. The earliest indication of trouble was when the first MICAS images were taken on November 6, 1998 (Fig 1(a)). These images showed that the frame was severely corrupted by stray light, clearly visible in the left portion of the frame, which has a background brightness level of about 1000 DN above the right portion. Of secondary importance, but still disconcerting, was the presence of several large smudges in the frame where the CCD has decreased sensitivity. It was quickly apparent from these images that the original autonav image processing software loaded onto the spacecraft would not work due the large variation in the background DN levels caused by the stray light. Additional images taken on November 17 further showed that the intensity of the stray light varied as a function of the angle between the camera boresight direction and the direction to the sun (hereafter referred to as the cone angle). Intensive efforts were then undertaken to characterize the problem by taking images at various cone angles. It was found that there were two regimes of stray light problems; between cone angles of 110 to 180 degrees, the light was in the form of a "card" as seen in the first frames which increased in

intensity as a function of increasing cone angle, and at cone angles less than 110 degrees, a "blowtorch" of light in the upper third of the frame appeared (Fig 1(b)). The former problem is more benign because the image is still usable at long exposure durations; the latter more serious because the blowtorch saturates very quickly and bleeds onto the remainder of the frame at exposures longer than about 10 seconds.

Another unanticipated problem surfaced in early December 1998 when a set of 27 images of a dense star cluster were taken and downlinked. The purpose of the images was to characterize the geometric distortions in the camera focal plane which, if not modelled properly, would cause errors in the orbit solutions. It was found that the standard model of the distortion field used by the flight software did not accurately represent the high frequency variations. In particular, the fitted parameters to the model computed from the locations of the star image centroids resulted in residual rms errors of nearly 1 pixel. If a beacon asteroid were one AU distant, a one pixel centroid error would amount to a spacecraft position fix error of about 2000 km – much larger than desirable. For comparison, the same model used on the Voyager and Galileo missions had rms errors of less than 0.1 pixel. Thus, two software fixes were needed; one to handle the stray light, another to account for the unusual distortions.

A third major problem with MICAS that affected autonav was the reduced sensitivity of the

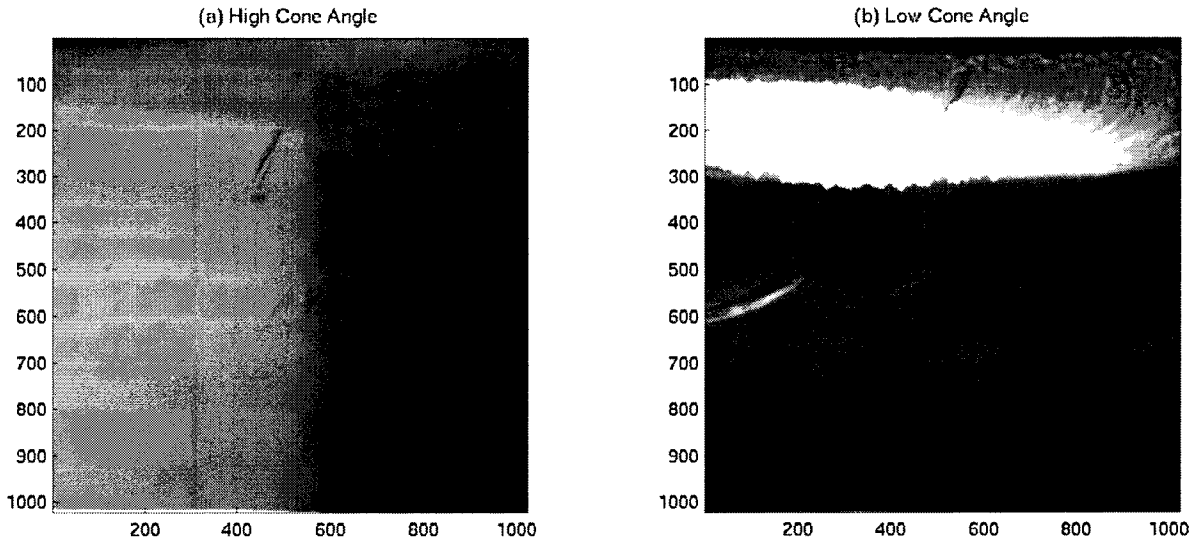

Fig. 1: images showing effects of stray light at two cone angles.

camera. Originally, the preflight specifications called for the ability to image stars and asteroids with visual magnitudes of 11.0 or brighter at reasonable exposure durations of tens of seconds. This allowed the use of up to 80 beacon asteroids, and provided a good geometric spread of beacons for any given OD opportunity which reduced the effects of asteroid ephemeris and other systematic, as well as random, errors. It also ensured that several stars would be in each frame for determining the absolute pointing direction. In actuality, only objects whose magnitudes were brighter than 9.0 – 9.5 were reliably detected. This severely constrained our choices for beacons; only the brighter asteroids were available, and the system had to cope often with frames which had fewer stars. As a consequence, the original charter of autonomy was descoped, with the autonav team on the ground taking responsibility

for carefully selecting the beacons which had the proper brightness and sufficient stars in the background.

For these and other reasons, the initial use of the computational elements of the autonav system, originally planned for early December, was delayed while new software was developed and tested (other autonav components which were not affected by the camera

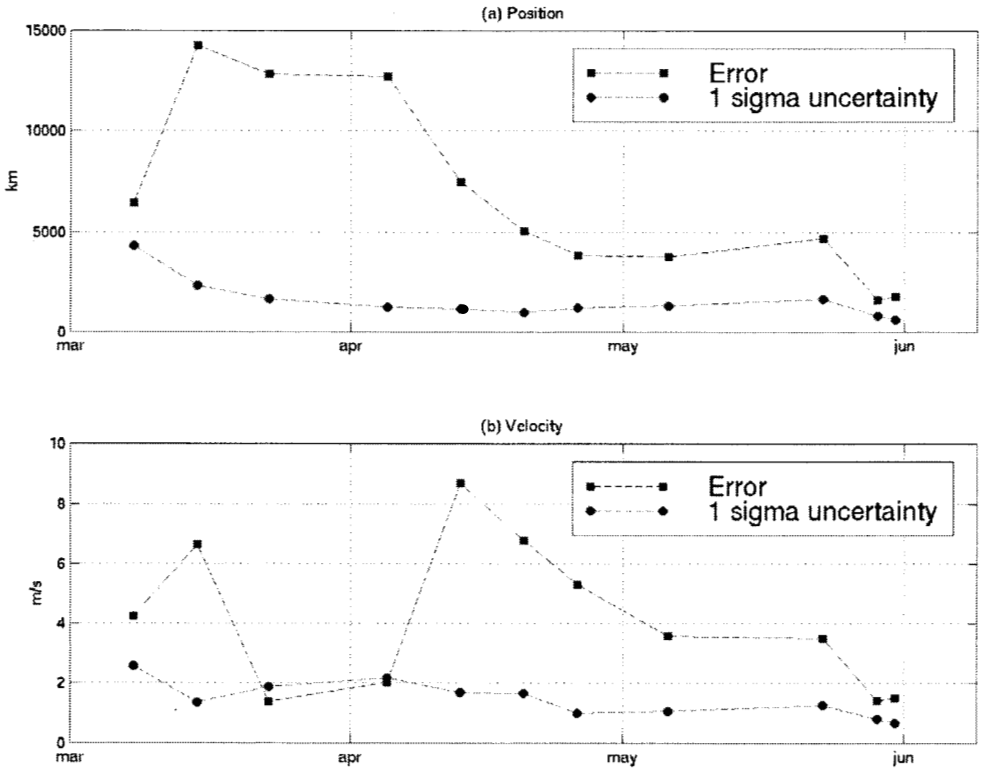

Fig. 2: position and velocity errors and associated uncertainties during the M4 time period.

problems were started shortly after launch and performed nominally). In the meantime, images were still being taken by the spacecraft and downlinked to the ground for analysis.

4.2 - OD Results During Interplanetary Cruise

During flight, two opportunities for uploading upgrades to the autonav software were available; the first was on February 8, 1999, and the second on June 10, 1999 (dubbed the "M4" and "M6" loads, respectively). The M4 load included new algorithms for initial registration and updated parameters for the original distortions model. M6 included the new distortion model and its associated parameters, the RSEN tracking code, and new code to perform frame differencing, whereby a "background" frame is first taken and differenced with the normal beacon frame. Differencing was added because the stray light was fairly stable over small differences in cone angle, and the differencing removed most of the large spatial wavelength variations from the stray light. For evaluating the performance of the onboard autonav, many of the raw images, as well as centroiding results, were downlinked for analysis. Also, since standard radio navigation was concurrently being performed on the

ground, the results from the onboard OD could be compared with the radio results (which is accurate to better than 100 km in heliocentric space).

The first onboard OD took place on February 18 using sightings of 4 asteroids, combined with the data from a "seed" file which had image centroids for pictures taken on January 7, 20, 26, and February 1, processed on the ground. After three more such campaigns on February 22, March 1, and March 8, the spacecraft computed its first completely autonomous onboard solution on March 8 using data starting on February 18. Subsequently, OD updates were performed at roughly weekly intervals with little ground intervention until early June,

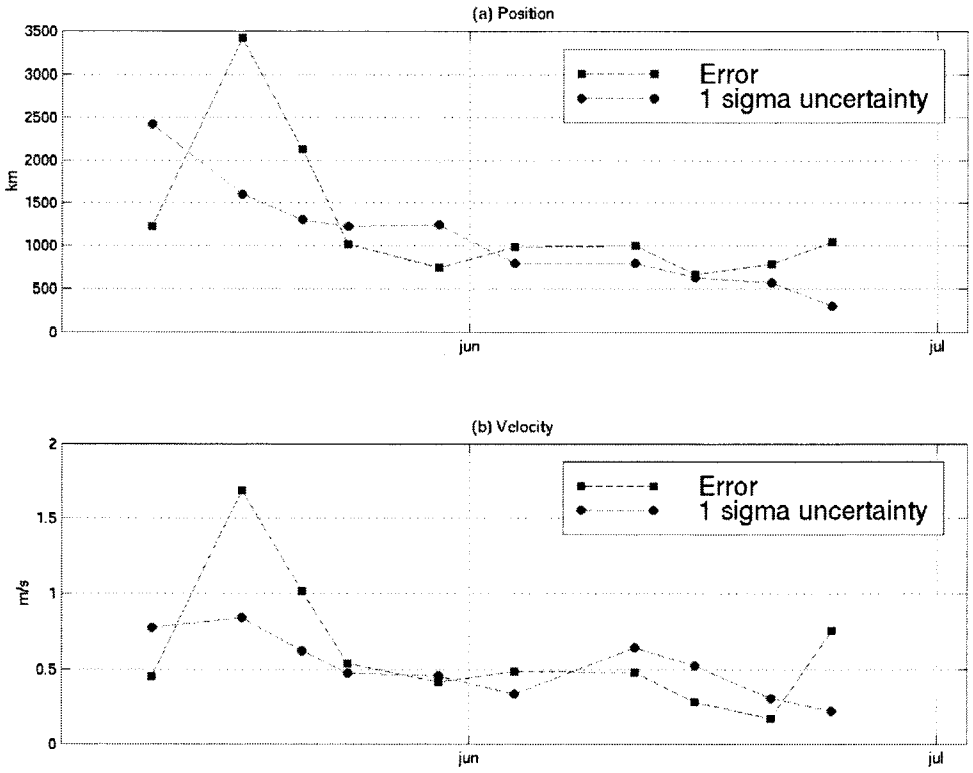

Fig. 3: position and velocity errors and uncertainties during theM6 time period.

1999, with a gap in mid-May due to testing of experimental non-navigation software components. Fig. 2 plots the onboard OD accuracies during the M4 time frame as the difference of the autonav estimated positions and velocities with the ground radio results. For simplicity, the rss of the three position and velocity components are used. Also plotted are the 1-sigma uncertainties in the position and velocity. The plot shows that after an initial position error of about 6000 km, the discrepancy ballooned to over 14,000 km before settling back down to the 4000 km range, and finally ending up less than 2000 km. The cause for the large jump was found to be an erroneous parameter describing the alignment of the camera; once this was fixed, it took several weeks before the data using the incorrect parameter exited the data arc. The marked improvement at the end was due to the addition of the bright asteroid Vesta to the solution which provided a larger angle between the observation directions than had been seen before (this can be seen in the position sigma, which dropped to about 600 km where before it had been hovering in the 1600 km range). The velocity errors varied between about 1.5 and 8.5 m/s. Both position and velocity errors were about an order of magnitude larger than preflight estimates, and usually several factors larger than the formal uncertainties.

This can be primarily ascribed to the loss of beacons and the lower fidelity distortion model, but another factor was the fact that the onboard estimates of RCS thruster firings provided by ACS were found to be off by nearly a factor of 2. The latter was corrected by M6.

Fig. 3 plots the onboard OD errors and associated sigmas computed after the M6 software load. Following the initial position error of around 1200 km, the results increased to 3400 km in mid-June before settling down to a 700-1000 km range. The sudden increase was, once again, due to erroneous parameters, which was corrected in mid-June. The steady state velocity errors during this time period was less than 0.5 m/s, and dropped as low as 0.2 m/s in parts of the arc. Clearly, the improvements in software and beacon selections had an effect on the OD accuracies, which were now only a factor of 2 or 3 larger than preflight estimates. Also note that with the M6 software, the errors are generally consistent with the formal uncertainties of the solutions. Overall, the system was now performing about as well as it could in theory given the geometrical constraints provided by the data. The loss of many beacons is reflected in the plots of the position uncertainties, which can vary markedly from one week to the next as opportune beacons rise and set. Had the camera response been as expected, the curve would have been much smoother and the solutions would not have been as sensitive to the gain or loss of any single beacon.

Between the upload of the corrected parameters in mid-June and the latter part of July, orbit determination onboard the spacecraft was fairly independent of ground intervention. Due to a combination of the reduced geometric information and unmodelled high frequency camera distortions, however, the accuracy being obtained was not sufficient to support fully autonomous maneuver planning. In particular, the maneuver planner was set to execute TCMs at Encounter (E) – 20, 10, 5, and 2 days, but only if the discrepancy between the predicted and desired target conditions was larger than 2 sigma of the formal error. A further complication was that the spacecraft performed a rehearsal of the encounter sequence in early July, which involved executing an actual maneuver based on simulated data for encounter. For these reasons, each TCM computed onboard was checked on the ground to see its effect on the spacecraft trajectory based on the more accurate radio data. Using this information, TCMs at E - 20 and E - 2 days were cancelled. The E – 10 day TCM was executed by carefully planning the rehearsal TCM to be in a direction which would be advantageous. The E – 5 day TCM computation was based on ground radio data due to an autonav flight software error which occurred on July 21 and caused corruption of the autonav file system.

4.3 – Approach Results

Based on ground data for the size, shape, and albedo of Braille, it was predicted that Braille would be visible in the MICAS CCD about 3 – 5 days before encounter. The actual first sighting of Braille occurred at E – 3 days, but only after extensive image processing on the ground. Following the TCM 5 days prior to encounter, the spacecraft had been taking multiple frames of Braille, but they were not used for OD because the onboard software was unable to detect any signal. Many of the frames were, however, sent to the ground for analysis. By registering the absolute orientation of the frames using a nearby guide star and then co-adding all the downlinked frames taken on July 26, a faint signal appeared which was likely due to Braille. Subsequent images taken and downlinked on July 27 verified that the signal, though still quite dim, was indeed from Braille. A ground computed OD solution based on these data indicated an ephemeris error for Braille of about 350 km (representing about 2 sigma change from the nominal ephemeris). This is graphically presented in Fig. 4 in

the B-plane coordinate system (the B-plane coordinate system is a plane centered on the target body and perpendicular to the incoming asymptote, with vertical and horizontal axes known as B•R and B•T, respectively). The current OD placed the spacecraft in the B-plane at a B•R of 248 km and B•T of 371 km, with an uncertainty of about 50 km (the ellipse in Fig. 4 represents the 1 sigma uncertainty). This was large enough to warrant a ground computation of a manuever at E – 1 day to retarget the spacecraft, which was executed nominally. The flyby aimpoint is located in this plane at a B•R of 12 km, and a B•T of -9 km.

Following the E – 1 day TCM, a set of 18 frames of Braille was taken. The onboard software was still unable to locate Braille in any of these images due to its signal being below the detectability threshold. From analyzing 5 of these images on the ground, the spacecraft's current flyby location was computed to be at a B•R of 7.6 km and B•T of 21 km, with an uncertainty of roughly 16 km (Fig. 4). Due to the size of the uncertainty, a planned E – 18 hr maneuver was cancelled.

At 11:30 UTC on July 28, 18 more frames of Braille were taken. Inspection of spacecraft telemetry data indicated that autonav had finally locked onto signal from Braille. Unfortunately, a latent bug in the autonav software caused a spacecraft safing event, wiping out the latest onboard trajectory result. Following recovery of the spacecraft from safing, three of the images were downlinked, and a ground OD solution computed. This solution showed that the spacecraft was now at a B•R of 4.0 km and a B•T of 6.0 km, with an uncertainty of around 11 km (Fig. 4). This result was somewhat disconcerting in that it represented a greater than 1 sigma deviation from the previous solution and furthermore, that the spacecraft was uncomfortably close to an impact trajectory. Thus, a TCM was developed to retarget back to the nominal aimpoint and uplinked to execute at E – 6 hours. The turnaround time from receiving the images to developing a maneuver sequence took no more than 1 hour.

Following the TCM, four sets of Braille images were taken at E –4, 3, 2, and 1.5 hours. These processed onboard normally, but did little to change the onboard estimate of the flyby trajectory, which, was computed to be at its targeted aimpoint location. This information was handed over to the RSEN subsystem at E – 27 minutes.

4.4 – RSEN Results

Real time Doppler data taken during encounter indicated that the flyby had proceeded safely and the spacecraft was operating normally. Following encounter, however, the downlinked close approach images did not show the asteroid, either in the CCD or APS frames, indicating an RSEN tracking failure. After all the onboard pictures were played back, only two frames showed Braille; both were in the CCD and taken approximately 15 minutes after closest approach. This, along with telemetry produced by RSEN and the ACS system, enabled the autonav team to reconstruct the events during encounter and pinpoint the probable cause of the failure.

Using all pre-encounter images of Braille and the two post-encounter ones, the actual flyby location was computed to be at a B•R of 25.8 km and B•T of –11.7 km, with an uncertainty of 1.5 km (Fig. 4). This was slightly greater than 1 sigma of its targeted location, but still well within the ability of RSEN to track. The telemetry from RSEN showed that APS frames between E – 27 minutes and E – 24 minutes were being processed normally, but that the

signal from Braille had not appeared above a preset threshold. At E – 24 minutes, a spurious signal appeared above the threshold and spoofed RSEN into making about a 30 km error in the B-plane estimate. This error biased subsequent commanded pointing directions to a level where Braille was no longer in the APS FOV. However, using ACS telemetry of pointing angles and the reconstructed trajectory, it was determined that Braille was still within the CCD FOV until E – 5 minutes, and then briefly again at E – 3 minutes. Unfortunately, due to onboard storage constraints, neither the science nor autonav teams scheduled CCD images to be taken and stored prior to E – 2 minutes. By E + 15 minutes, the spacecraft had reverted to its pre-RSEN trajectory, and captured Braille in the CCD. In addition, four APS images taken during this time also had Braille in its FOV, but only in one of them was there a signal of any note. This signal was barely above the background, and could only be identified because it matched the predicted location of Braille.

The question then remained as to why the signal from Braille was so weak. Data informally presented by various sources prior to encounter had led the autonav team to believe that Braille would produce a strong response in the APS with the selected exposure durations (about 5 seconds), and the threshold values for detection were set accordingly. The lower than expected signal in the post-encounter CCD frames, and the lack of signal in the APS, however, indicated that Braille was perhaps 10 times dimmer than predicted. Also, the unusual shape of Braille probably presented an aspect to the approaching spacecraft which decreased its integrated brightness. Finally, the response of the APS detector was nonlinear and poorly characterized, resulting in an optimistic expectation of its sensitivity.

Such factors led to autonav being unable to detect Braille in the E – 27 to E – 24 minute period. The noise spike which perturbed RSEN's estimate of the trajectory indicates a design flaw in the software, which did not check for persistence of a signal before accepting it as real. This feature was not thought necessary on the assumption that the signal from the target would be so strong that the effects of noise would be negligible. This was verified in an experiment where the two post-encounter CCD images of Braille were used as data for RSEN. RSEN was able to pinpoint the flyby to within two km using just the two images, and the noise in these frames had no effect. However, even had RSEN had not been spoofed by the noise, the result would have been the same because the trajectory knowledge prior to initiating RSEN was not sufficient to track the asteroid open loop beyond the E – 2.5 minute time period.

5 - SUMMARY AND CONCLUSIONS

After the initial difficulties with parameters following the M4 and M6 uploads, autonav, at least as far as image processing and OD are concerned, was performing its task more or less without ground intervention. The maneuver computation routines were slightly less autonomous in that the state information it was given occasionally had to be uplinked from the ground, but it still worked onboard as designed. For the interplanetary cruise portion of the mission, the autonav system was deemed validated by the project, and indeed, following the Braille encounter, had been in control of the spacecraft with little intervention and no radio backup until the star tracker failure in November. The accuracies obtained by the M6 version of the software are more than sufficient for this mission phase, and a planned M7 load would have improved on its overall reliability. During the approach phase to a target, however, the ground would have to intervene because the OD accuracy is not good enough to

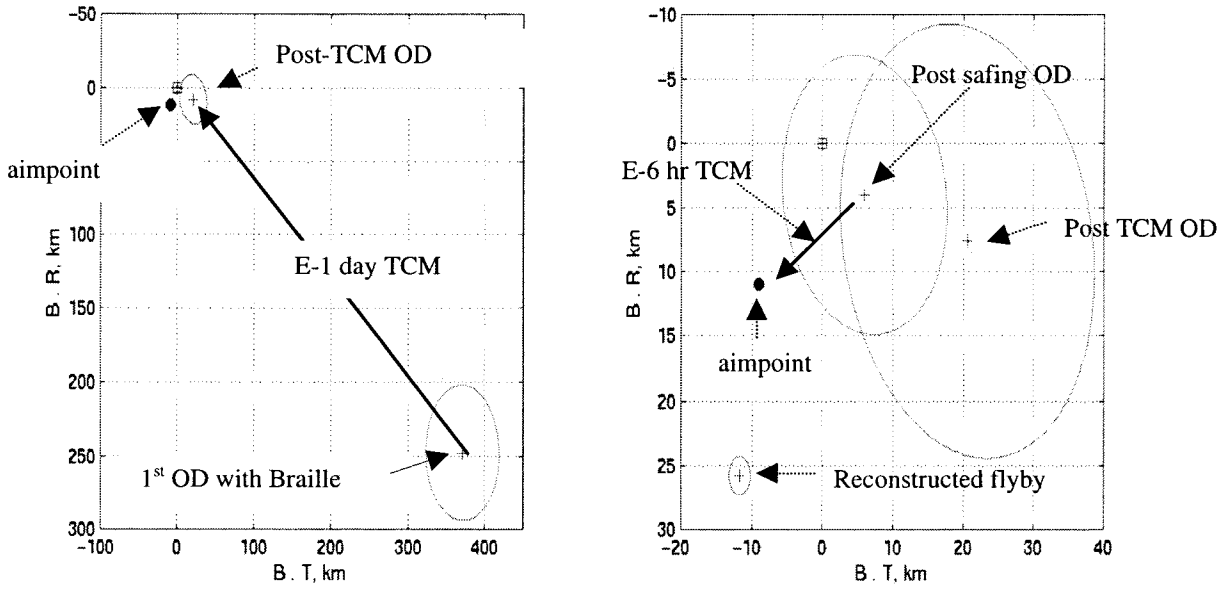

Fig. 4: B-plane at encounter.

provide a robust sequence of maneuvers which would take the spacecraft progressively closer to its aimpoint. This failure, however, can be attributed almost entirely to the fact that the MICAS camera did not meet its design requirements. For future missions, it is clear from the DS1 experiment that, with a well designed and characterized camera, an autonav system could easily perform both mission phases reliably.

Following the lessons learned from the Braille encounter, the RSEN tracking code has been modified to make it more robust for the Borrelly encounter. These modifications, however, cannot be used as a substitute for proper characterization of the flyby target. This is an important point to note for other missions which rely on an autonomous system for closed loop tracking of a small, relatively unknown object, such as the STARDUST mission which will use a very similar algorithm during its comet encounter. In this matter, DS1 has performed a valuable role in real-world testing of unproven techniques which can enhance the scientific returns from more conventional missions.

ACKNOWLEDGEMENT

The work described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

REFERENCES:

[Bhas 97] S. Bhaskaran, J. E. Riedel, S. P. Synnott, "Autonomous Nucleus Tracking for Comet/Asteroid Encounters: The STARDUST Example", *Paper AAS 97-628, AAS/AIAA Astrodynamics Conference*, Sun Valley, ID, August 1997.

[Bhas 98] S. Bhaskaran, S. D. Desai, P. J. Dumont, B. M. Kennedy, G. W. Null, W. M. Owen, J. E. Riedel, S. P. Synnott, R. A. Werner, "Orbit Determination Performance Evaluation of the Deep Space 1 Autonomous Navigation System", *Paper AAS 98-193, AAS/AIAA Spaceflight Mechanics Meeting*, Monterrey, CA, February 1998.

[Desa 97] S. D. Desai, S. Bhaskaran, W. E. Bollman, C. A. Hallsell, J. E. Riedel, S. P. Synnott, "The DS-1 Autonomous Navigation System: Autonomous Control of Low-Thrust Propulsion Systems", *AIAA Paper 97-38819, AIAA Guidance, Navigation and Control Conference*, New Orleans, LA, August 1997.

[Lieb 67] P. B. Liebelt, *An Introduction to Optimal Estimation*, Addison Wesley, 1967.

[Raym 96] M. D. Rayman, D. H. Lehman, "NASA's First New Millennium Deep Space Technology Validation Flight", *IAA Paper IAA-L-0502, Proceedings of the Second IAA International Conference on Low-Cost Planetary Missions*, Laurel, MD, April 1996.

[Ried 97] J. E. Riedel, S. Bhaskaran, S. P. Synnott, S. D. Desai, W. E. Bollman, P. J. Dumont, C. A. Hallsell, D. Han, B. M. Kennedy, G. W. Null, W. M. Owen, R. A. Werner, B. G. Williams, "Navigation for the New Millennium: Autonomous Navigation for Deep Space 1", *Proceedings of the 12th International Symposium on Space Flight Dynamics*, Darmstadt, Germany, June 1997.

[Vaug 92] R.M. Vaughan, J. E. Riedel, R. P. Davis, W. M. Owen, S. P. Synnott, "Optical Navigation for the Galileo Gaspra Encounter", *AIAA Paper 92-4522, AIAA/AAS Astrodynamics Conference*, Hilton head, S. C., August 1992.